

# PinchLens: Applying Spatial Magnification and Adaptive Control-Display Gain for Precise Selection in Virtual Reality

Fengyuan Zhu\*

Department of Computer Science, University of Toronto

Mauricio Sousa‡

Department of Computer Science, University of Toronto

Ludwig Sidenmark†

Department of Computer Science, University of Toronto

Tovi Grossman§

Department of Computer Science, University of Toronto

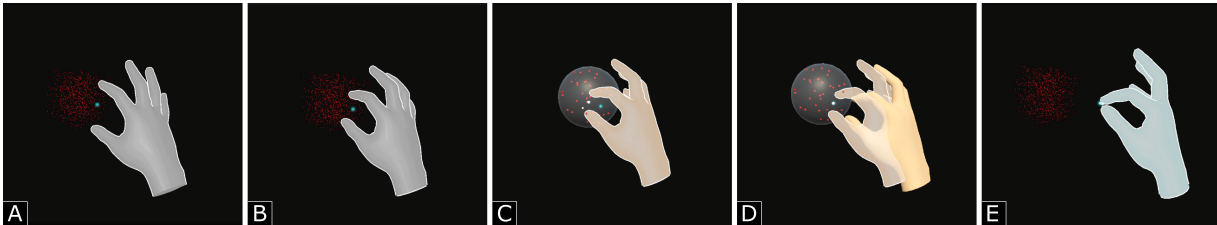


Figure 1: The PinchLens technique overview: A) The user aims to select a small target in a dense cluster. B) the user begins a pinching gesture. C) A magnifying bubble expands the scale of nearby targets when the user reaches a Semi-Pinch. D) An adaptive control-display gain that dampens the motion of the user’s hand for precision. E) A final pinch selects the target.

## ABSTRACT

We present PinchLens, a new free-hand target selection technique for acquiring small and dense targets in Virtual Reality. Traditional pinch-based selection does not allow people to precisely manipulate small and dense objects effectively due to tracking and perceptual inaccuracies. Our approach combines spatial magnification, an adaptive control-display gain, and visual feedback to improve selection accuracy. When a user starts the pinching selection process, a magnifying bubble expands the scale of nearby targets, an adaptive control-to-display ratio is applied to the user’s hand for precision, and a cursor is displayed at the estimated pinch point for enhanced visual feedback. We performed a user study to compare our technique to traditional pinch selection and several variations to isolate the impact of each of the technique’s features. The results showed that PinchLens significantly outperformed traditional pinch selection, reducing error rates from 18.9% to 1.9%. Furthermore, we found that magnification was the dominant feature to produce this improvement, while the adaptive control-display gain and visual cursor of pinch were also helpful in several conditions.

**Index Terms:** Human-centered computing—Human computer interaction (HCI)—Interaction techniques

## 1 INTRODUCTION

Current hand-tracking techniques for Virtual Reality (VR) allow people to use freehand gestures to interact with virtual objects naturally, without the need for handheld controllers. These techniques facilitate interactions with virtual objects, emulating the way people interact with the physical world, offering dexterity and a natural way to manipulate virtual objects. These advancements bridge the gap between the digital and physical realms, paving the way for a future where technology seamlessly integrates with human intuition.

However, one of the principal challenges facing current VR applications is the difficulty in manipulating small objects within densely packed environments. There are many situations where such interactions may be needed. For example, in a VR LEGO® application, a user may need to acquire and assemble tiny LEGO® pieces<sup>1</sup>. In an electronics training app, a user may want to rewire connections on a densely packed breadboard. Or, in a 3D jewelry app, a designer may want to choose a specific seed bead to add to a bracelet (Fig. 2). In the physical realm, such dexterous interactions are enabled by the high-fidelity visual and haptic sensory information we receive from the real world. However, selecting small, densely-packed objects in VR remains a daunting task [43, 50, 70]. This predicament can be attributed to several factors. First, current tracking techniques have not yet achieved the millimeter-level accuracy necessary to support highly dexterous operations [50, 66, 67]. Second, there is limited dexterity due to limitations in motor control causing jittering during mid-air movements [53]. Finally, human perceptual errors and visual acuity issues associated with stereo rendering further complicate the task [15, 39, 44, 74]. These challenges compromise the effectiveness of virtual interactions, thereby falling short of the fluidity observed in interactions within the physical world.

Various methods have been proposed to facilitate target selection with small, densely-packed objects in VR applications. Some techniques that have been explored include separating the degrees of operations [55], slicing the selection area [56], and utilizing multiple rays [22, 49]. However, these methods are largely indirect and deviate from the paradigm of direct interaction that could potentially leverage real-world skills. Recent studies [50, 59] have emphasized the need for precise selection with direct manipulation, although the applications in VR environments remain limited.

In this work, we contribute PinchLens, a novel technique for selecting small-scale targets in dense VR environments, integrating spatial magnification and adaptive control-display gain, while retaining an intuitive pinch-based selection paradigm (Fig. 1A). PinchLens is inspired by facilitation techniques used in 2D desktop interactions, including the Bubble Cursor [33], Bubble Lens [57], and Adaptive Control Display gain [14]. Upon initiation of the pinch selection process (triggered by a “Semi-Pinch” gesture), a virtual bubble appears, magnifying the target region and aiding in visually acquiring

\*e-mail: fy Zhu@dgp.toronto.edu

†e-mail: lsidenmark@dgp.toronto.edu

‡e-mail: mauricio@dgp.toronto.edu

§e-mail: tovi@dgp.toronto.edu

<sup>1</sup>LEGO®: <https://www.lego.com>




Figure 2: Many physical tasks require precise interaction with tiny dense objects, which are still difficult to enable in Virtual Reality.

the target (Fig. 1B). Simultaneously, the user’s virtual hand movement is visually slowed down using a control-gain ratio for precise movement (Fig. 1C). A selection event is triggered when the pinch is complete (Fig. 1D), at which point the magnification and motion dampening are smoothly transitioned back to normal.

We conducted a user study to assess the performance of several variations of our technique, in comparison to traditional pinch selection, to identify the effects of the individual design features of PinchLens. The results suggest that PinchLens reduces the error rate from 18.9% to 1.9%, and decreases the overall distance between the pinch point and the target by 43%, shortening the completion time under dense conditions by 18%, without affecting the completion time under sparse conditions. Furthermore, we found that magnification was the primary factor for improvement, while the adaptive control-display gain and the visual cursor could also enhance performance under certain conditions. We conclude by discussing the implications of our findings and avenues of future work.

## 2 RELATED WORK

The PinchLens technique draws on a variety areas of prior research, including target selection techniques for desktop environments, strategies developed specifically for VR-based target selection, and adaptive control-display gain techniques in VR.

### 2.1 Desktop Target Selection Techniques

Human-Computer Interaction researchers have developed various techniques that support target selection. Target distance and width have been shown to have a significant impact on selection, where small and distant targets are more difficult to select [4, 51]. A typical approach to address challenging selections is to adjust the target width or distance to reduce the precision required for correct selection [4]. Previous work has achieved effective selection performance by expanding targets [6, 27, 52, 54, 57, 81], area cursors [33, 40, 75, 80] and using adaptive control-display gain [14, 18, 76]. Specifically, we adapt concepts from the and Bubble Cursor [33], which dynamically adjusts its size to increase the effective width of nearby targets, and the Bubble Lens [57], which magnifies nearby targets, activated automatically through kinematic triggering, making them larger in both visual and motor-space.

### 2.2 Target Selection Techniques in Virtual Reality

Selection is an essential category of interaction in VR [7, 47], and is used to indicate the intention of manipulating and interacting with objects in the environment. Previous work identified two main 3D selection techniques – ray-casting [24, 34, 46, 61], which enables distant interaction by casting a virtual ray, and using virtual representations of the user’s hand (3D cursor) [25, 38, 48, 62]. The virtual hand metaphor has the advantage over ray-casting in that targets are easily manipulated after selection, while ray-casting allows easier selection of distant targets [62]. In this work, we focus on the virtual hand metaphor for selection, but we also take inspiration from previous work on ray-casting to select small targets in dense environments.

Ray-casting utilizes a ray starting from the user’s hand or controller, and a selection is performed by intersecting the ray with an object. Ray-casting techniques have proven effective for selection [45]. However, small targets can be challenging to select due to motor limitations [5]. Furthermore, selection becomes difficult in dense environments due to occlusion, which reduces the effective target size, causing users to carefully manipulate the position and orientation of the ray to select occluded targets [68]. Previous work has proposed disambiguation techniques to improve the selection of small and occluded targets by adding additional steps to the interaction [5, 68], progressive refinement from a set of selection candidates [43], or predictive models of selection [21, 65]. Although these techniques improve selection, their interaction methods utilize an indirect manipulation approach, diverging from how people interact in the physical world.

The virtual hand metaphor tries to mimic hand interaction with the physical world [19, 37, 69, 71]. However, fine-grained selection of small targets in dense environments can be difficult to perform with direct manipulation due to limitations in tracking and human motor control [11, 32]. As such, researchers have developed techniques to support direct selection and direct manipulation of small targets in dense environments. Ma et al. [50] leverage the user’s selection history and the distance between objects and the hand to determine the interacted object. Frees et al. [26] leverage a dynamic control-display gain to allow more precise movements in the display space. Finally, Periverzov et al. [59] leverage hand motions to infer the intended target to enable the selection of small targets. We take inspiration from these works by leveraging the spatial relationship between hand and objects, combined with spatial magnification and dynamic control-display gain. Our study shows that our technique provides enhanced selection performance with millimeter-scale precision in dense environments.

### 2.3 Adaptive Control-Display Gain in Virtual Reality

Balakrishnan highlighted that while pointing in the physical world is constrained by physical laws, pointing in the virtual world does not have to abide by the same constraints [4]. The concept of adaptive control-display gain has shown to be helpful for interaction by increasing precision and has guided various techniques in VR. Early research, such as the family of Go-Go techniques [9, 62, 63], implement nonlinear mapping of hand movements, manipulating virtual hands to effectively reach remote objects beyond the physical reach of the user. Expanding on this idea, Frees et al. [26] used dynamic control-display gain to alter the movement of controlled targets in relation to different speeds of virtual hands, providing varying levels of sensitivity for target manipulation in VR. Adaptive control-display gain has also been extensively used to aid in passive haptic retargeting by modifying the position of virtual hands for redirected touch [2, 35, 41, 42], or haptic retargeting with graspable objects [3, 28, 29, 31]. These body warping techniques have also been used recently to predict user intention and facilitate virtual target acquisition within arm’s range [23, 30, 77, 78].

These works have focused on employing adaptive control-display




Figure 3: The ‘‘Semi-Pinch’’ state and the triggering process. The  $t_1$  and  $t_2$  should be set to different values to avoid false positives.

gain in VR to reach and manipulate objects of easily graspable target sizes. Limited research explores whether adaptive control-display gain can help improve the precision of selecting small targets in a dense environment. This gap prompts us to incorporate similar techniques to modify user hand motions, to support precise operations over small and densely situated targets.

### 3 PINCHLENS TECHNIQUE

PinchLens is designed to improve the performance of target selection for small and dense targets in VR, utilizing a direct manipulation input metaphor. The technique was developed to maintain as much of the physical action of real-world pinching as possible, while leveraging target facilitation mechanisms to improve its precision. This section provides a detailed account of the PinchLens implementation, encompassing the following primary components: 1) a Semi-Pinch gesture to trigger the facilitation process. 2) a spherical bubble lens that magnifies nearby targets in visual and motor space, 3) an adaptive control-display gain that enables precise hand movements, with an accompanying recovery process to realign the virtual and physical hands, and 4) a selection cursor to visualize the pinch location. In addition, we describe a global input data stabilizer to reduce the jitter of the hand and fingers resulting from the VR motion tracking. The pseudocode for each component is provided in the supplementary material, and the example code can be accessed in the GitHub Repository<sup>2</sup>

#### 3.1 Semi-Pinch Trigger Mechanism

The ‘‘Pinch’’ [8, 12] has become a widely adopted gesture [58, 60] for natural selection in VR, akin to the mouse click in desktop environments. Not only does the pinch gesture provide a sense of direct manipulation, but it also provides an explicit and tactile moment of selection, when the thumb and index figure come in contact [73]. With PinchLens, the pinch still confirms the selection, but we introduce an intermediate state, prior to the pinch occurring, to activate the mechanisms that facilitate the selection process.

We leverage a Semi-Pinch gesture to indicate the user’s intent to select, and to trigger the PinchLens’s selection facilitation features. When the tips of the index finger and thumb approach each other (less than a threshold,  $t_1$ ) before making contact, our system enters the Semi-Pinch state. Conversely, the user can exit this state without performing a selection by moving the index and thumb tips apart and beyond a threshold  $t_2$ . Fig. 3 shows PinchLens’s transition between states. Through pilot testing, we selected the following threshold values that we gave users reasonable control:  $t_1 = 3.5cm$  and  $t_2 = 4.0cm$ . Once the semi-pinch occurs, the midpoint between the thumb tip and index finger tip define the pinch point. The proximity of nearby targets are also considered when triggering the PinchLens, to avoid unnecessary false activations. Specifically, if the nearest target is greater than 0.2m from the pinch location, we opt not to activate the PinchLens technique.

#### 3.2 Spatial Magnification

Inspired by desktop target selection techniques such as the Bubble Cursor [33] and Bubble Lens [57], we adapted a spherical expansion

<sup>2</sup>Example code repository: <https://github.com/zhufyaxel/PinchLens>

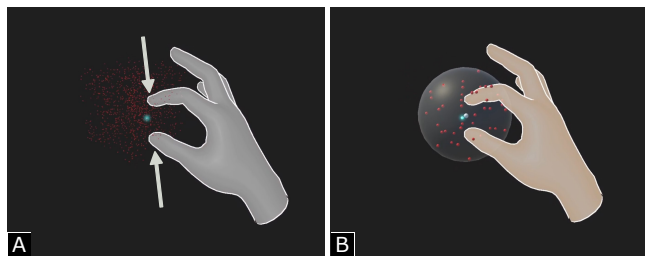


Figure 4: The process of Spatial Magnification. When the Semi-Pinch is triggered, the bubble will appear and magnify nearby objects while hiding targets outside the bubble.

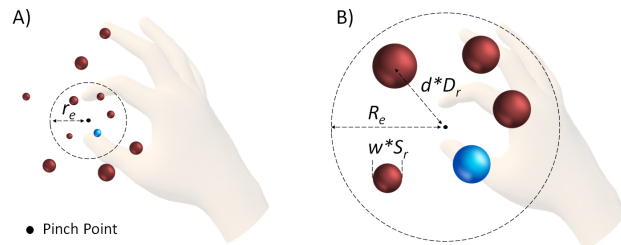


Figure 5: The magnification method and parameters. a) Targets within the threshold distance,  $r_e = 1.67cm$ , are magnified. b) The expanded sphere has an outer distance of  $R_e = 5.0cm$ . The width,  $w$ , of targets are magnified by a factor of  $S_r = 4x$ , and their distance,  $d$ , from the pinch point are magnified by a factor of  $D_r = 3x$ . The blue dots represents the target to be selected.

lens for VR. This lens is triggered when the semi-pinch occurs. The mechanism uses the ‘‘pinch point’’ as the selection center to magnify nearby targets in visual and motor space. Here we describe the details of the expansion lens, which is illustrated in Fig. 4.

##### 3.2.1 Defining the Magnification Area

A spatial 3D environment is prone to occlusions. Therefore, spatially magnified targets may cause occlusion and difficulty in identifying selectable targets [56]. Inspired by previous works [56, 64], we implemented a ‘‘target visibility’’ mechanism during the selection process where targets outside the magnification area are rendered transparent to minimize occlusion. The following steps define the target visibility mechanism: when the Semi-Pinch is activated, targets further away than a threshold ( $r_e$ ) from the Semi-Pinch point are rendered invisible. Targets within the threshold are spatially magnified (Fig. 4B). The invisible targets are restored when the Semi-Pinch gesture is canceled or the selection is completed. This strategy effectively minimizes visual occlusion caused by magnification while decreasing the chance that users mistakenly select an unintended target.

##### 3.2.2 Magnification Method

Previous work on target magnification has shown that separately magnifying target distances and sizes is an effective technique to efficiently use the space within the magnification lens while retaining spatial relationships [16, 57]. A uniform magnification ratio for size and distance may provide inadequate magnification if the magnification is contained in a small magnification area, or exacerbate occlusions in the environment. We define separate ratios for magnifying distance and size, and use the pinch point as the center of our magnification. When the semi-pinch occurs, all targets within the distance,  $r_e$ , to the pinch point are magnified (Fig. 5A). These targets are expanded outward along the sphere’s radial direction to a

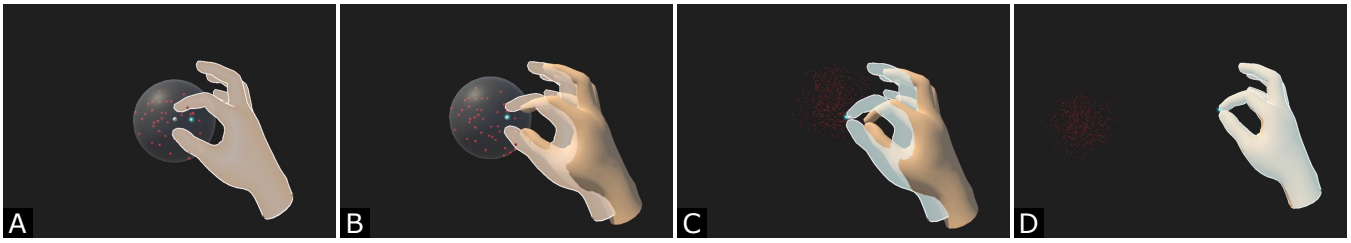


Figure 6: The Motion Dampening and Recovery Process. A) The Semi-Pinch is triggered, invoking the magnifying bubble; B) Once the user moves their hands for aiming, an adaptive control-display gain will be applied to make the visual hand (transparent) moves slower than the physical hand (opaque, which won't show in real application); C) After pinching, the motion dampening ends D) The virtual hand then redirects back to the physical position.

maximum distance of  $R_e$ , in proportion to a distance ratio  $D_r$ , and a scale ratio,  $S_r$ , to increase the target size (Fig. 5B).

After several iterations and internal pilot studies, we set the distance ratio  $D_r$  to 3.0, and the size expansion ratio,  $S_r$  to 4.0. The radius of the area of interest was set at  $r_e = 1.67cm$ , and therefore the radius of the expanded sphere was  $R_e = 5.0cm$ . This approach allows us to retain the spatial relationships of targets within the lens while maximizing target sizes within the lens area. Please refer to the pseudocode (Algorithm 1) provided in the supplementary material for a full description of the expansion process.

### 3.3 Adaptive Control Display Gain

To improve the selection accuracy when the spatial magnification is triggered, we implemented an adaptive control-display gain mechanism that decelerates the user's virtual hand movements to enhance precision. The process encompassing this adaptive gain and the subsequent recovery phase is depicted in Fig. 6. Previous techniques have established methods to adjust the speed of virtual hands within VR while keeping users unaware of the manipulation [3, 79]. However, many of these techniques rely on preexisting knowledge of the target location [17, 30] and typically apply these adjustments throughout the entire range of arm movement. Our approach differs in that we do not presume prior knowledge of the target location and only adapt the users' movements in a limited range much shorter than the full arm's reach.

Our approach to the *adaptive control-display gain* can be likened to the modification of a mouse cursor's sensitivity: when the user initiates a Semi-Pinch, our technique dampens the movement of the virtual hand, reducing the distance traversed in the virtual space relative to the real physical movement. After the selection is completed or interrupted, we initiate a "recovery" phase, where we gradually synchronize the visual hands with the user's actual physical hand position. We use the same approach as in previous work, which adds extra movement to the virtual hand toward the physical hand position [29, 30]. The movement speed is set in proportion to the real hand's movement speed, ensuring consistency of visual and motor cues. The adaptive movements are facilitated by adjusting the root speed of the virtual hand, followed by corresponding updates to the local skeleton data. Where the expansion technique is clearly noticeable to the user, the adaptive control-display gain technique can easily go unnoticed to the user.

After several iterations, we established a control-display ratio of  $\alpha = 0.5$  once the PinchLens is triggered. This ratio effectively reduces the virtual hand movement speed by half, improving the precision of selection within the spatially magnified environment. The recovery speed ratio is set to  $\beta = 0.3$  for natural recovery. This means that the virtual hand will initiate an extra movement towards the physical hands, and this additional velocity will be 30% of the absolute velocity relative to the physical hands. As a result, the total velocity of the virtual hands becomes the sum of the velocity of the physical hands and this additional velocity component. This




Figure 7: Visual Cursor for Pinch. After we trigger the semi-pinch gesture, a small spatial cursor will appear and show where the estimated pinch position

condition holds until the virtual hand aligned with the physical hand. For a detailed representation of our adaptive control-display gain algorithm, please refer to the pseudocode Algorithm 2 provided in the supplementary material.

### 3.4 Visual Pinch Cursor

Although conventional widgets such as ray-casting [49] or pinch proxy<sup>3</sup> offer ways to visualize the pointer location in VR, their design does not support direct manipulation or precise interaction with nearby objects. For the PinchLens, we implemented a visual cursor to visualize the pinch point (midpoint between the index finger tip and the thumb tip) once the semi-pinch occurs (Fig. 7). The white sphere-shaped cursor indicates the 3D position of the pinch point, enabling users to better comprehend the spatial relationship between their pinch point and interacted targets. To avoid visual clutter when the user is not performing selections, the visual cursor is only visible during the Semi-Pinch state. We established the cursor's diameter to be 0.5cm after our internal pilot testing. This size ensures the cursor's visibility and enables precise selection, without causing significant occlusion to surrounding targets.

### 3.5 Global Stabilizer

Previous work has suggested that current computer vision-based hand tracking approaches commonly used for hand tracking and interaction have inherent errors and noise that negatively affect target acquisition [66, 67]. Even in ideal tracking conditions, jittering is a common issue due to the natural tremor of the hands, especially during precise movements [72]. Because the pinch point and visual cursor are positioned in between the fingertips, sudden changes in the tracked finger positions may change the pinch point and make precise selection of small targets difficult.

Our global stabilizer is inspired by Casiez et al.'s 1€ filter [13] that utilizes a low-pass filter to refine noisy input signals for high precision and responsiveness. We adopt a similar strategy to stabilize

<sup>3</sup>Quest Pinch Pose: <https://ocul.us/3Ve8J2i>

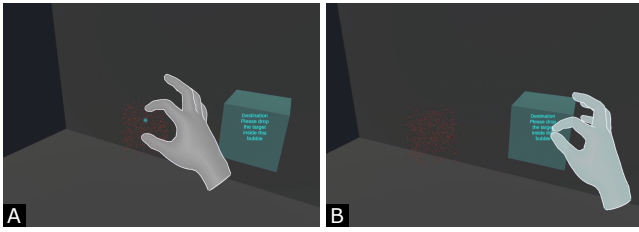


Figure 8: Evaluation environment. We instructed the participants to use a pinch gesture to select A) a target (highlighted in blue) from the dense cluster of small objects, and B) transfer into the blue cube.

the inaccuracy of the hand data for target selection. The  $1\epsilon$  filter uses a low-pass filter to refine noisy input signals for high precision and responsiveness and has been proven suitable for stabilizing hand data in VR [20]. For our global stabilizer, we apply a filter over the entire hand skeleton model data, as users may rely on the virtual representation of the whole hand for interaction. To avoid unexpected distortion of the skinned mesh [36], we first apply the  $1\epsilon$  filter on the root pose and then apply another  $1\epsilon$  filter on the rotational skeleton data in its local space. This combination of filters ensures that the hand moves more naturally. After filtering, we reapply the skinned mesh on the filtered positions. The pinch point is then calculated based on the filtered data. Filtering is applied at all times, irrespective of the pinch state. Please refer to Algorithm 3 in the supplementary material for the pseudocode.

## 4 EXPERIMENT

We conducted a user study aimed at evaluating the performance of PinchLens in the selection process of small targets within dense environments. Our objective was to measure the overall effectiveness of PinchLens and gain insight into the contribution of each individual component to its functionality. To accomplish this, we conducted a comprehensive performance comparison with multiple variations of our technique in comparison to a baseline condition.

In our initial implementation of the study, the baseline condition was designed as a basic pinch selection technique. The technique did not include any visual enhancements to mimic the real world scenario, and the pinch point needed to intersect a target for a selection to occur. However, our pilot studies indicated that the absence of any visual feedback and the requirement to intersect the targets resulted in a nearly 100% error rate. Previous research [22, 49] has proposed that visual hints could effectively reduce selection errors and could be integrated into the baseline. Consequently, we integrated a “highlight effect” that highlights the nearest target when a Semi-Pinch is triggered, and relaxed the selection criteria to be the nearest target. This modification created a more equitable baseline condition for our study. We refer to this modified condition as the *enhanced baseline*.

Our technique itself consists of three major components: the magnifying bubble, the adaptive control-display gain approach, and the visual pinch cursor. To balance the effective conditions and study length, we implemented three variations of the PinchLens technique to test the performance from each component separately. As such, the conditions included in our study were as follows: 1) Enhanced Baseline Pinch, 2) Pinch Cursor (excluding adaptive control-display gain and spatial magnification), 3) Pinch Motion (incorporating adaptive control-display gain but excluding spatial magnification), and 4) PinchLens (including all enhancements). In all conditions, the closest target to the pinch point was highlighted and selected. To ensure fairness, the global stabilizer was applied to all conditions, including Enhanced Baseline Pinch.

### 4.1 Task

Previous studies in selecting targets from dense clusters of objects adopted a “click and go next” approach [33, 49]. However, we opted for a “pick and place” loop that better represents the whole manipulation process in VR [9, 10]. In this study, we instructed participants to use a pinch gesture to select a specific goal target from the dense cluster, move the target while maintaining the pinch, and then place it into a semi-transparent cube located 20cm to the right of the cluster. Fig. 8 shows the evaluation environment. The goal target in the cluster was highlighted in blue, and the instruction “Please drop the target inside this Cube” was displayed on the cube for placing the targets. Participants were allowed to load the next trial once they had placed the target within the cube area. The cluster of targets was positioned in front of the participant at a distance of 30cm. The cluster and cube heights were adjusted for each participant to ensure they could comfortably reach them.

Drawing from previous research [15, 39] and through iterative development, we established the following parameters for target sizes and densities within our study: 0.75mm and 1.5mm for the target sizes in diameters, and density levels of 500 and 1000 targets randomly distributed within an 8cm \* 8cm \* 8cm cube volume. These study conditions created the types of environments which are known to be particularly problematic for existing VR selection techniques, and were selected as we were specifically interested in the selection of millimeter-scale targets.

In summary, the study used the following independent variables and levels:

- **Technique:** Baseline Pinch, Pinch Cursor, Pinch Motion, PinchLens.
- **Target Size:** 0.75mm (small), 1.5mm (large).
- **Target Density:** 500 targets (sparse), 1000 targets (dense).

### 4.2 Apparatus

In this study, we used a Meta Quest 2<sup>4</sup> and used its built-in hand-tracking capability (version 41) for our study. Previous work has shown that Meta Quest 2 hand-tracking (version 33) has an average fingertip positional error of 1.1cm and an average finger joint angle error of 9.6° [1]. We built the evaluation prototype in Unity3D 2021.3.4f1<sup>5</sup> in Windows 11 operating system, and ran the software in a Laptop with Nvidia RTX 3070 Graphic Card<sup>6</sup>. To provide a more stable fit with the Quest 2 during the study, we used a Quest Elite 2 Strap<sup>7</sup> for stability and comfort.

### 4.3 Participants

We recruited 12 participants aged 19 to 35 (4 female, average age = 23.6 years, sd = 3.4). All participants reported being right-handed and not visually impaired. 5 out of 12 participants reported previous familiarity with VR and the remaining participants reported having limited or no experience. The study was approved by our institutional ethics board and participants provided informed consent prior to commencing the study.

### 4.4 Procedure

At the beginning of the study, we carefully adjusted the interpupillary distance and headset position for each participant until they confirmed optimal vision clarity within the virtual environment. To verify this, we rendered a slim line of random phrases (0.5cm width per character), positioned 35cm away from the participant, to test their vision accuracy in VR.

<sup>4</sup>Meta Quest 2: <https://www.meta.com/quest/products/quest-2/>

<sup>5</sup>Unity Archive: <https://bit.ly/3Cw8E1C>

<sup>6</sup>Laptop Spec: <https://Inv.gy/3CQHbJo>

<sup>7</sup>Elite Strap: <https://bit.ly/3ECAFFn>

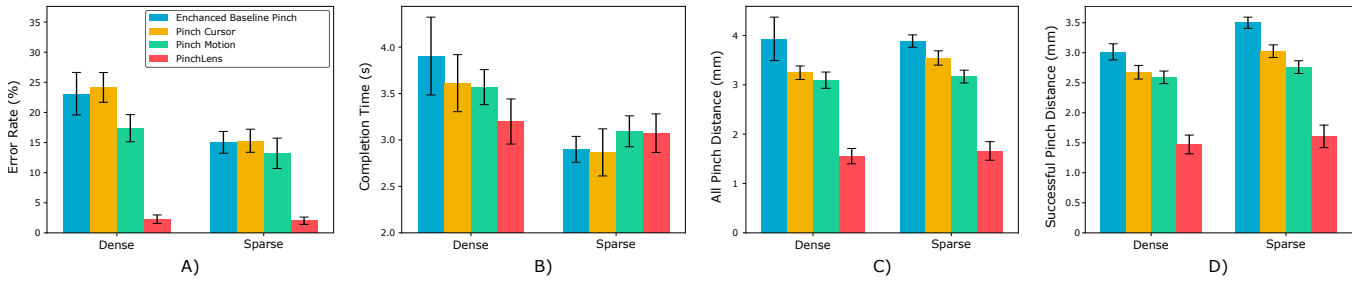


Figure 9: Evaluation results showing A) Error Rate (%), B) Completion Time (s), C) All Pinch Distance (mm), and D) Pinch Distance (mm) from successful trials. Error bars represent the standard error of the means.

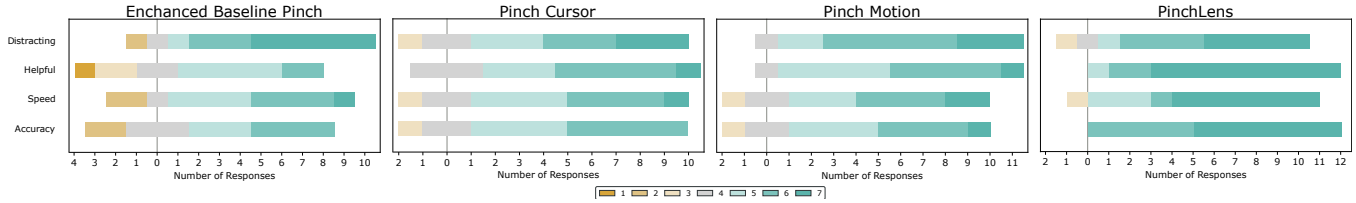


Figure 10: Results from the 7-point Likert scale questionnaire. For clarity in visual representation, scores pertaining to distractions were inverted: “0” represented a very distracted state, while “7” indicated no distractions.

Next, we generated the target cluster and adjusted its height until the participant could comfortably reach it. The cluster was loaded with our large and sparse density conditions. We then introduced the participants to the concepts of Semi-Pinch and pinch selection, with only the pinch cursor and highlight as the supportive features. Once participants gained sufficient Semi-Pinch proficiency and were able to perform pinch selections comfortably, the study formally commenced.

We used a Balanced Latin Square to counterbalance the technique order and randomized the order of Size  $\times$  Density in each technique. The participants were asked to perform 16 pick and place trials for each combination. When each new technique was encountered, there was a short warm-up session for the participant to get familiar with the technique. For each trial, one goal target, rendered blue, was randomly chosen from the target cluster. All other targets were rendered red, so that participants could easily identify the goal target. Participants had to pinch the cube to start each trial. Each trial allowed for a maximum of five attempts before the trial was considered unsuccessful and skipped. If participants reached the maximum failure count in a trial, the semi-transparent cube would change color and display the text “Maximum Failures to this Trial, pinch the cube to continue.” At this moment, pinching the cube resumed the study by loading the next trial. At the end of the study, to gather subjective feedback, we asked the participants to fill out a 7-point Likert scale questionnaire, and we conducted a semi-structured interview. The study duration was approximately 45 minutes. Between each technique, participants rested for a minimum of one minute. Participants were also told they could rest at any time. In total, each participant performed  $4 \times 2 \times 2 \times 16 = 256$  trials.

In order to measure the task performance, we measured the *error rate*, the *completion time*, and the *pinch distance*. The error rate was the proportion of trials resulting in the user selecting an incorrect target, calculated by dividing the number of erroneous trials by the total number of selections. A trial was marked as “successful” if the correct target was selected at the first attempt, while any incorrect selection would mark the trial as an error. Completion time refers to the time it took to complete the trial, calculated as the time from when the cube was pinched to start the trial, to when the target was released in the cube. Finally, the pinch distance was the Euclidean distance between the goal target center and the pinch point. We

performed separate analyses on the pinch distance for successful selections and all selections.

## 4.5 Results

We recorded data in CSV format. We used *statsmodels* and *pandas* in Python to process and analyze the data. Unless otherwise stated, our analysis was performed with a three-way repeated measures analysis of variance (RM-ANOVA) with Technique, Density, and Target Size as independent variables. Holm-Bonferroni-corrected post-hoc tests were used when applicable. A summary of the results is provided in Table 1.

### 4.5.1 Error Rates

An RM-ANOVA showed that the technique significantly affected the error rate ( $F(3, 33) = 23.305, p < .001$ ). Target density also showed a statistical significance ( $F(1, 11) = 9.401, p < .05$ ). Fig. 9A shows that the PinchLens technique significantly improved the overall accuracy, with  $2.08\% \pm 0.70\%$  for dense conditions and  $1.82\% \pm 0.60\%$  for sparse, compared to the Enhanced Baseline Pinch condition that had  $22.92 \pm 3.53\%$  in dense conditions and  $14.84 \pm 1.81\%$  in sparse conditions. The post-hoc pairwise comparisons revealed significant differences in error rate between PinchLens and other techniques in both dense and sparse conditions ( $p < .001$ ). Furthermore, the post-hoc pairwise comparisons showed that density significantly affected the error rate for the pinch cursor and pinch motion techniques ( $p < 0.05$ ). We found no significant interaction effects. These results show that density increases the error rate. Nevertheless, the results suggest that our approach, in particular the expansion lens, substantially improves accuracy in both the sparse and dense conditions.

### 4.5.2 Completion Time

Fig. 9B shows the completion time across techniques and densities. For this part of the analysis, we excluded trials that reached the maximum amount of selection failures. There was a main effect of density ( $F(1, 11) = 21.671, p < .001$ ), and a significant two-way interaction between technique and density ( $F(3, 33) = 4.066, p < .05$ ). In the dense condition, pairwise comparisons showed that PinchLens was significantly faster than Baseline Pinch ( $p < .01$ ) and Pinch Motion ( $p < .05$ ). No other significant differences were found in dense or sparse conditions. These results indicate that when

Techniques	Density	Error-Rate(%)	Completion Time(s)	Successful Pinch Distance(mm)	All Pinch Distance(mm)
Enhanced Baseline Pinch	Dense	22.92±3.53	3.90±0.42	3.01±0.13	3.93±0.44
Pinch Cursor	Dense	23.96±2.47	3.61±0.31	2.68±0.11	3.24±0.14
Pinch Motion	Dense	17.18±2.26	3.57±0.19	2.59±0.11	3.10±0.17
PinchLens	Dense	2.08±0.70	3.19±0.24	1.48±0.16	1.56±0.16
Enhanced Baseline Pinch	Sparse	14.84±1.81	2.90±0.14	3.50±0.09	3.89±0.13
Pinch Cursor	Sparse	15.10±1.91	2.87±0.25	3.03±0.11	3.54±0.14
Pinch Motion	Sparse	17.19±2.26	3.09±0.17	2.76±0.11	3.17±0.13
PinchLens	Sparse	1.82±0.60	3.07±0.20	1.61±0.19	1.67±0.19

Table 1: Summary of results across techniques and target densities.

the target is in a dense environment, the PinchLens technique could reduce the completion time (3.19±0.24 seconds) compared to the Baseline Pinch condition (3.90±0.42 seconds), without changing the completion time in the sparse condition.

#### 4.5.3 Pinch Distance

For pinch distance of all selections, the RM-ANOVA showed that only the technique has a significant impact ( $F(3, 33) = 28.656, p < .001$ ). We found no significant interactions. Fig. 9C shows these results. A significant difference was found between all the techniques to the Baseline Pinch: PinchLens ( $p < .001$ ), Pinch Cursor ( $p < .001$ ), and Pinch Motion ( $p < .001$ ). Similarly, PinchLens was significantly different from both the Pinch Cursor ( $p < .001$ ) and Pinch Motion ( $p < .001$ ). However, no statistically significant difference was detected between the Pinch Cursor and Pinch Motion.

For the pinch distance of only successful trials, the RM-ANOVA showed significant main effects for technique ( $F(3, 33) = 54.514, p < .001$ ) and density ( $F(1, 11) = 39.953, p < 0.01$ ). There was also a significant technique  $\times$  density ( $F(1, 11) = 39.953, p < 0.01$ ) two-way interaction. We show the results in Fig. 9D. Post-hoc comparison of the two-way interaction showed that for dense targets, the PinchLens had significantly less pinch distance over all other techniques (all  $p < .001$ ), and the Baseline Pinch had a significantly higher pinch distance compared to Pinch Cursor ( $p < .05$ ) and Pinch Motion ( $p < .001$ ). There was no significant difference between Pinch Cursor and Pinch Motion. After post-hoc comparisons, PinchLens was found to be the most effective technique for sparse targets ( $p < .05$  between the Pinch Cursor and Pinch Motion, and  $p < .001$  between the Pinch Cursor and other two conditions).

PinchLens significantly reduced the pinch distance compared to other methods. Furthermore, we observed that the pinch cursor and pinch motion could also reduce the successful pinch distances and all pinch distances compared to the baseline.

#### 4.5.4 Subjective Feedback

We asked participants to complete a 7-point Likert scale questionnaire for each technique. Fig. 10 shows the results of the questionnaire. The results suggest that the PinchLens received overall positive feedback compared to the other conditions.

During the interview, we were interested to know whether the participants could perceive the adaptive control-display gain when applied. Four out of 12 participants reported that they perceived the adaptive control-display gain with the PinchLens techniques. Three out of 12 participants also reported perceiving the adaptive control display gain without the provided magnifying bubble, but only one out of 12 participants reported it both with PinchLens and the condition with no magnifying bubble. The general feedback from the post-interview was positive. Our participants found the PinchLens “was helpful” (P1), “improved confidence” (P3), provided “strong willing for use” (P4), “reduced mental load” (P8), “easy and accurate” (P9), “quicker to use” (P10). Other comments included that the technique needs extra effort to learn and adapt (P3, P4, P9) and that

it “potentially could be designed as a movable object with proper strategy” (P12).

## 5 DISCUSSION AND FUTURE WORK

The PinchLens technique performs statistically and subjectively well, with significantly lower error rates (from  $\sim 20\%$  to  $\sim 2\%$ ) and pinch distance (from  $\sim 3.9\text{cm}$  to  $\sim 1.6\text{cm}$ ). These improvements are significant in enabling fine-grained selection with few added interaction steps for the user, while maintaining a direct manipulation interaction metaphor. Additionally, we found that PinchLens led to shorter completion times by  $\sim 18\%$  under dense conditions. Finally, we found that the pinch cursor and Adaptive Control Display Gain approaches alone can improve performance compared to baseline pinching. Together, these findings reveal the vital role of spatial magnification in selecting small targets in dense environments and bolstering the performance of the PinchLens technique, while also highlighting the benefits of our other components.

Our study employed a strategy that compared the variances of PinchLens techniques at different levels of component combination. This approach was formulated to balance the number of effective conditions with the study’s duration, mindful of the potential physical fatigue participants might face from the repetitive pick-place loop in our task. Though we recognize that investigating additional conditions could yield further insights, we deemed the current design optimal for our investigation. Future research may delve into these areas, carefully weighing their potential impact.

Regarding target density, we found that PinchLens is proficient in sparse and dense conditions. It is also worth noting that even our sparsest “sparse” condition combined with our “large” size is denser than that seen in recent prior studies [22, 50]. Our more challenging conditions, combined with results comparable to those from these studies and techniques, showcase the utility of PinchLens for selection in dense environments. We also found that the target size did not affect PinchLens’s performance. Although we only compared two sizes (0.75mm and 1.5mm), PinchLens’s high performance gives us confidence in its ability to select small targets. That said, it would be important in the future to test the PinchLens in conditions when targets are much larger or sparsely distributed. The technique was designed to only activate when a small target is nearby, so we would expect that its performance would be similar to traditional baseline pinching when selection is “easy”, but future studies are needed to ensure this is the case.

Our study used a pick-and-place loop rather than a simple selection task as in previous works. We made this decision to enable the subsequent manipulation and release of the target after selection, a common scenario in hand-based interaction [10]. By adopting this strategy, we foster a more comprehensive user experience to emulate the entire operation loop rather than a selection process without subsequent steps. However, our study focused on evaluating selection and we did not evaluate PinchLens combined with subsequent manipulation and release, which can be a valuable future research direction. In particular, in some tasks, like placing a Lego block, the precision of the placement of a small object will be just

as important as the precision of the initial selection of that object. We believe the concepts PinchLens used to facilitate selection could also be applied to facilitate placement, but future research needs to explore the proper trigger for the facilitation mechanisms.

We utilized the visual presence of targets in a straightforward solution, where red dots were randomly distributed within a cubic space to symbolize potential targets, with a distinct blue dot serving as the target expected to be selected. The magnification was uniformly applied within a spherical area, ensuring the target remained visible from all viewing angles within the designated area of interest. Given the importance of visual cues in VR for target acquisition, further evaluations could be conducted to explore the potential influence of different visual factors.

For the PinchLens technique, the Semi-Pinch acts as a trigger for the spatial magnification and adaptive control-display gain. This trigger is powerful as it gives users control over the timing and position of the added features in one gesture. It also enables users to easily cancel the interaction by reverting the hand to an idle state. We found in our study that the simple threshold-based implementation of Semi-Pinch was adequate for interaction. Future studies could investigate further capabilities of the Semi-Pinch, such as continuous Semi-Pinching or more granular interaction states.

Compared to previous Bubble Lens techniques [57], our trigger method employs gesture detection as a state (Semi-Pinch) rather than kinematic triggering based on overall speed detection. We posit that gesture-based triggering could reduce the likelihood of false positives and ensure the smoothness of the selection process. Future explorations may also consider evaluating dynamic parameters and triggering mechanisms similar to those found in the Bubble Lens to determine potential enhancements.

We applied the PinchLens technique for targets within arm's reach. However, our three primary components: spatial magnification, adaptive control-display gain, and the pinch cursor, do not need to be confined to arm's length and can be used for interactions from afar with proper remapping. This insight suggests that PinchLens may be effective in selecting distant targets and targets within arm's reach. Future research could consider integrating our method with arm extensions or ray-casting techniques. Evaluating these enhancements to determine their potential in magnifying spatial objects and applying adaptive control-display gain for precise operations is a promising avenue for future studies.

During our study, we utilized fixed parameters for pinch detection, the rate of magnification, and the size of the cursor. All participants could adapt to these parameters after training and completing the tasks. An additional potential future direction could be to dynamically offer adjusted parameters during the pinch-selection loop. Such designs could involve the implementation of an adaptive magnification ratio [57] and a dynamically resized spatial cursor similar to the bubble cursor technique [33]. These refinements could be devised and assessed in subsequent studies.

Finally, our study took place in a controlled abstract environment, allowing us to study the techniques across various target size and density conditions. We envision numerous VR applications in which the PinchLens technique could be particularly useful. As depicted in (Figure 2), there are many real-world applications where precise selection of small objects is needed, such as when working with electronics, jewellery, or toy models. Such applications could be ported to VR systems for either entertainment or training purposes, and utilize PinchLens when precise selection is required. Implementations and evaluations of such scenarios are left to future work.

## 6 CONCLUSION

Supporting target selection with small and dense objects in VR is an essential task. However, due to tracking and perceptual errors, it is difficult for users to select small-scale targets with direct manipulations like a pinch. In this paper, we tackled the issue by providing

a novel target selection approach named PinchLens based on 2D desktop facilitation techniques. The technique maintains a direct manipulation interaction metaphor, and combines spatial magnification, a pinch cursor, and adaptive control-display gain to improve the precision of target selection with small-scale objects in dense environments. Evaluation results showed that our techniques significantly improved target selection performance with small-scale objects, with a lower error rate and pinch distance, and significantly reduced completion time within the dense environment. Results also showed that each component of PinchLens has its unique value under different conditions. Our work shows that spatial magnification and adaptive control-display gain can play an important role in enabling the selection and manipulation of millimeter-scale targets in dense environments, and can extend the capabilities of free-hand interaction in VR.

## REFERENCES

- [1] D. Abdulkarim, M. Di Luca, P. Aves, M. Maaroufi, S.-H. Yeo, R. C. Miall, P. Holland, and J. M. Galea. A Methodological Framework to Assess the Accuracy of Virtual Reality Hand-Tracking Systems: A Case Study With the Meta Quest 2. *Behavior Research Methods*, pp. 1–12, 2023. doi: 10.3758/s13428-022-02051-8
- [2] P. Abtahi and S. Follmer. Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2018. doi: 10.1145/3173574.3173724
- [3] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In *Proceedings of the 2016 Chi Conference on Human Factors in Computing Systems*, pp. 1968–1979, 2016. doi: 10.1145/2858036.2858226
- [4] R. Balakrishnan. “Beating” Fitts’ Law: Virtual Enhancements for Pointing Facilitation. *International Journal of Human-Computer Studies*, 61(6):857–874, 2004. doi: 10.1016/j.ijhcs.2004.09.002
- [5] M. Baloup, T. Pietrzak, and G. Casiez. Raycursor: A 3d Pointing Facilitation Technique Based on Raycasting. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019. doi: 10.1145/3290605.3300331
- [6] B. B. Bederson. Fisheye Menus. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*, pp. 217–225, 2000. doi: 10.1145/354401.354782
- [7] D. Bowman, E. Kruijff, J. J. LaViola Jr, and I. P. Poupyrev. *3D User Interfaces: Theory and Practice, CourseSmart ETextbook*. Addison-Wesley, 2004.
- [8] D. Bowman, C. Wingrave, J. Campbell, and V. Ly. Using Pinch Gloves (tm) for Both Natural and Abstract Interaction Techniques in Virtual Environments. 2001. doi: 10.1145/3229434.3229470
- [9] D. A. Bowman and L. F. Hodges. An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, pp. 35–ff, 1997. doi: 10.1145/253284.253301
- [10] D. A. Bowman and L. F. Hodges. Formalizing the Design, Evaluation, and Application of Interaction Techniques for Immersive Virtual Environments. *Journal of Visual Languages & Computing*, 10(1):37–53, 1999. doi: 10.1145/323663.323667
- [11] D. A. Bowman, D. B. Johnson, and L. F. Hodges. Testbed Evaluation of Virtual Environment Interaction Techniques. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, pp. 26–33, 1999. doi: 10.1145/323663.323667
- [12] D. A. Bowman and C. A. Wingrave. Design and Evaluation of Menu Systems for Immersive Virtual Environments. In *Proceedings IEEE Virtual Reality 2001*, pp. 149–156. IEEE, IEEE, 2001. doi: 10.5555/580521.835855
- [13] G. Casiez, N. Roussel, and D. Vogel. 1€ Filter: A Simple Speed-Based Low-Pass Filter for Noisy Input in Interactive Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2527–2530, 2012. doi: 10.1145/2207676.2208639
- [14] G. Casiez, D. Vogel, R. Balakrishnan, and A. Cockburn. The Impact of Control-Display Gain on User Performance in Pointing Tasks. *Human-*



- Computer Interaction*, 23(3):215–250, 2008. doi: 10.1145/2485760.2485792
- [15] M. Chessa, G. Maiello, A. Borsari, and P. J. Bex. The Perceptual Quality of the Oculus Rift for Immersive Virtual Reality. *Human-Computer Interaction*, 34(1):51–82, 2019. doi: 10.1109/VRW55335.2022.00029
- [16] M. Choi, D. Sakamoto, and T. Ono. Bubble Gaze Cursor + Bubble Gaze Lens: Applying Area Cursor Technique to Eye-Gaze Interface. In *ACM Symposium on Eye Tracking Research and Applications*, ETRA '20 Full Papers. Association for Computing Machinery, New York, NY, USA, 2020. doi: 10.1145/3379155.3391322
- [17] A. Clarence, J. Knibbe, M. Cordeil, and M. Wybrow. Unscripted Retargeting: Reach Prediction for Haptic Retargeting in Virtual Reality. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 150–159. IEEE, IEEE, 2021. doi: 10.1109/VR50410.2021.00036
- [18] A. Cockburn and A. Firth. Improving the Acquisition of Small Targets. In *People and Computers XVII—Designing for Society*, pp. 181–196. Springer, 2004. doi: 10.1007/978-1-4471-3754-2\_11
- [19] B. R. De Araújo, G. Casiez, and J. A. Jorge. Mockup Builder: Direct 3D Modeling on and Above the Surface in a Continuous Interaction Space. In *Proceedings of Graphics Interface 2012*, pp. 173–180. 2012. doi: 10.5555/2305276.2305305
- [20] M. C. de Castro, P. Joao, P. F. Rosa, and J. C. de Oliveira. Interaction by Hand-Tracking in a Virtual Reality Environment. In *2021 22nd IEEE International Conference on Industrial Technology (ICIT)*, vol. 1, pp. 895–900. IEEE, IEEE, 2021. doi: 10.1109/QoMEX51781.2021.9465407
- [21] G. De Haan, M. Koutek, and F. H. Post. IntenSelect: Using Dynamic Object Rating for Assisting 3D Object Selection. In *Ipt/egve*, pp. 201–209. Citeseer, Citeseer, 2005. doi: 10.5555/2385984.2386013
- [22] W. Delamare, M. Daniel, and K. Hasan. MultiFingerBubble: A 3D Bubble Cursor Variation for Dense Environments. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pp. 1–6, 2022. doi: 10.1145/3491101.3519692
- [23] S. Esmacili, B. Benda, and E. D. Ragan. Detection of Scaled Hand Interactions in Virtual Reality: The Effects of Motion Direction and Task Complexity. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 453–462. IEEE, IEEE, 2020. doi: 10.1145/3489849.3489866
- [24] P. M. Fitts. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology*, 47(6):381, 1954.
- [25] A. Forsberg, K. Herndon, and R. Zeleznik. Aperture Based Selection for Immersive Virtual Environments. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, pp. 95–96, 1996. doi: 10.1145/237091.237105
- [26] S. Frees, G. D. Kessler, and E. Kay. PRISM Interaction for Enhancing Control in Immersive Virtual Environments. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 14(1):2–es, 2007. doi: 10.1145/1229855.1229857
- [27] G. W. Furnas. Generalized Fisheye Views. *Acm Sigchi Bulletin*, 17(4):16–23, 1986. doi: 10.1145/22627.22342
- [28] B. Geslain, S. Besga, F. Lebrun, and G. Bailly. Generalizing the Hand Redirection Function in Virtual Reality. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces*, pp. 1–9, 2022. doi: 10.1145/3531073.3531100
- [29] E. J. Gonzalez, P. Abtahi, and S. Follmer. Reach+ Extending the Reachability of Encountered-Type Haptics Devices Through Dynamic Redirection in Vr. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 236–248, 2020. doi: 10.1145/3379337.3415870
- [30] E. J. Gonzalez, E. D. Chase, P. Kotipalli, and S. Follmer. A Model Predictive Control Approach for Reach Redirection in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2022. doi: 10.1145/3491102.3501907
- [31] E. J. Gonzalez and S. Follmer. Investigating the Detection of Bimanual Haptic Retargeting in Virtual Reality. In *25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–5, 2019. doi: 10.1145/3359996.3364248
- [32] T. Grossman and R. Balakrishnan. Pointing at Trivariate Targets in 3D Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 447–454, 2004. doi: 10.1145/985692.985749
- [33] T. Grossman and R. Balakrishnan. The Bubble Cursor: Enhancing Target Acquisition by Dynamic Resizing of the Cursor’s Activation Area. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 281–290, 2005. doi: 10.1145/1054972.1055012
- [34] T. Grossman and R. Balakrishnan. The Design and Evaluation of Selection Techniques for 3D Volumetric Displays. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, pp. 3–12, 2006. doi: 10.1145/1166253.1166257
- [35] D. T. Han, M. Suhail, and E. D. Ragan. Evaluating Remapped Physical Reach for Hand Interactions With Passive Haptics in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1467–1476, 2018.
- [36] S. Han, B. Liu, R. Cabezas, C. D. Twigg, P. Zhang, J. Petkau, T.-H. Yu, C.-J. Tai, M. Akbay, Z. Wang, et al. MEgATrack: Monochrome Egocentric Articulated Hand-Tracking for Virtual Reality. *ACM Transactions on Graphics (ToG)*, 39(4):87–1, 2020. doi: 10.1145/3386569.3392452
- [37] O. Hilliges, S. Izadi, A. D. Wilson, S. Hodges, A. Garcia-Mendoza, and A. Butz. Interactions in the Air: Adding Further Depth to Interactive Tabletops. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*, pp. 139–148, 2009. doi: 10.1145/1622176.1622203
- [38] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. A Survey of Design Issues in Spatial Input. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology*, pp. 213–222, 1994. doi: 10.1145/192426.192501
- [39] X. Jin, J. Meneely, and N.-K. Park. Virtual Reality Versus Real-World Space: Comparing Perceptions of Brightness, Glare, Spaciousness, and Visual Acuity. *Journal of Interior Design*, 47(2):31–50, 2022. doi: 10.1111/joid.12209
- [40] P. Kabbash and W. A. Buxton. The “prince” Technique: Fitts’ Law and Selection Using Area Cursors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 273–279, 1995. doi: 10.1145/223904.223939
- [41] L. Kohli. Redirected Touching: Warping Space to Remap Passive Haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 129–130. IEEE, IEEE, 2010. doi: 10.5555/2013876.2013909
- [42] L. Kohli, M. C. Whitton, and F. P. Brooks. Redirected Touching: Training and Adaptation in Warped Virtual Spaces. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 79–86. IEEE, IEEE, 2013. doi: 10.1145/2964284.2964293
- [43] R. Kopper, F. Bacim, and D. A. Bowman. Rapid and Accurate 3D Selection by Progressive Refinement. In *2011 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 67–74. IEEE, IEEE, 2011. doi: 10.1109/3DUI.2011.5759219
- [44] E. Kruijff, J. E. Swan, and S. Feiner. Perceptual Issues in Augmented Reality Revisited. In *2010 IEEE International Symposium on Mixed and Augmented Reality*, pp. 3–12. IEEE, IEEE, 2010. doi: 10.1145/2522848.2522885
- [45] D. M. Lane, S. C. Peres, A. Sandor, and H. A. Napier. A Process for Anticipating and Executing Icon Selection in Graphical User Interfaces. *International Journal of Human-Computer Interaction*, 19(2):241–252, 2005. doi: 10.1207/s15327590ijhc1902\_5
- [46] J. J. LaViola. Double Exponential Smoothing: An Alternative to Kalman Filter-Based Predictive Tracking. In *Proceedings of the Workshop on Virtual Environments 2003*, pp. 199–206, 2003. doi: 10.1145/769953.769976
- [47] J. J. LaViola Jr, E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, 2017.
- [48] J. Liang and M. Green. JDCAD: A Highly Interactive 3D Modeling System. *Computers & Graphics*, 18(4):499–506, 1994. doi: 10.1145/3139131.3139158
- [49] Y. Lu, C. Yu, and Y. Shi. Investigating Bubble Mechanism for Ray-Casting to Improve 3D Target Acquisition in Virtual Reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp.

- 35–43. IEEE, IEEE, 2020. doi: 10.1145/3604383.3604395
- [50] J. Ma, J. Qian, T. Zhou, and J. Huang. FocalPoint: Adaptive Direct Manipulation for Selecting Small 3D Virtual Objects. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(1):1–26, 2023. doi: 10.1111/cgf.12466
- [51] I. S. MacKenzie. Fitts’ Law As a Research and Design Tool in Human-Computer Interaction. *Human-Computer Interaction*, 7(1):91–139, 1992. doi: 10.1207/s15327051hci070\_3
- [52] J. D. Mackinlay, G. G. Robertson, and S. K. Card. The Perspective Wall: Detail and Context Smoothly Integrated. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 173–176, 1991. doi: 10.1145/108844.108870
- [53] S. Mayer, V. Schwind, R. Schweigert, and N. Henze. The Effect of Offset Correction and Cursor on Mid-Air Pointing in Real and Virtual Environments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2018. doi: 10.1145/3173574.3174227
- [54] M. McGuffin and R. Balakrishnan. Acquisition of Expanding Targets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 57–64, 2002. doi: 10.1145/503376.503388
- [55] D. Mendes, F. Relvas, A. Ferreira, and J. Jorge. The Benefits of Dof Separation in Mid-Air 3d Object Manipulation. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*, pp. 261–268, 2016. doi: 10.1145/2993369.2993396
- [56] R. A. Montano-Murillo, C. Nguyen, R. H. Kazi, S. Subramanian, S. DiVerdi, and D. Martinez-Plascencia. Slicing-Volume: Hybrid 3d/2d Multi-Target Selection Technique for Dense Virtual Environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 53–62. IEEE, IEEE, 2020.
- [57] M. E. Mott and J. O. Wobbrock. Beating the Bubble: Using Kinematic Triggering in the Bubble Lens for Acquiring Small, Dense Targets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 733–742, 2014. doi: 10.1145/2556288.2557410
- [58] A. K. Mutasim, A. U. Batmaz, and W. Stuerzlinger. Pinch, Click, or Dwell: Comparing Different Selection Techniques for Eye-Gaze-Based Pointing in Virtual Reality. In *ACM Symposium on Eye Tracking Research and Applications*, pp. 1–7, 2021. doi: 10.1145/3448018.3457998
- [59] F. Periverzov and H. İlieş. IDS: The Intent Driven Selection Method for Natural User Interfaces. In *2015 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 121–128. IEEE, IEEE, 2015. doi: 10.1109/3DUI.2015.7131736
- [60] K. Pfeuffer, B. Mayer, D. Mardanbegi, and H. Gellersen. Gaze+ Pinch Interaction in Virtual Reality. In *Proceedings of the 5th Symposium on Spatial User Interaction*, pp. 99–108, 2017. doi: 10.1145/3131277.3132180
- [61] J. S. Pierce, A. S. Forsberg, M. J. Conway, S. Hong, R. C. Zeleznik, and M. R. Mine. Image Plane Interaction Techniques in 3D Immersive Environments. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, pp. 39–ff, 1997. doi: 10.1145/253284.253303
- [62] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa. The Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, pp. 79–80, 1996. doi: 10.1145/237091.237102
- [63] I. Poupyrev, T. Ichikawa, S. Weghorst, and M. Billinghurst. Egocentric Object Manipulation in Virtual Environments: Empirical Evaluation of Interaction Techniques. In *Computer Graphics Forum*, vol. 17, pp. 41–52. Wiley Online Library, Wiley Online Library, 1998. doi: 10.1145/3485279.3485296
- [64] J. L. Reisman, P. L. Davidson, and J. Y. Han. A Screen-Space Formulation for 2D and 3D Direct Manipulation. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*, pp. 69–78, 2009. doi: 10.1145/1622176.1622190
- [65] G. Schmidt, Y. Baillot, D. G. Brown, E. B. Tomlin, and J. E. Swan. Toward Disambiguating Multiple Selections for Frustum-Based Pointing. In *3D User Interfaces (3DUI’06)*, pp. 87–94. IEEE, IEEE, 2006. doi: 10.1109/VR.2006.133
- [66] D. Schneider, V. Biener, A. Otte, T. Gesslein, P. Gagel, C. Campos, K. Čopič Pucihar, M. Kljun, E. Ofek, M. Pahud, et al. Accuracy Evaluation of Touch Tasks in Commodity Virtual and Augmented Reality Head-Mounted Displays. In *Symposium on Spatial User Interaction*, pp. 1–11, 2021. doi: 10.1145/3485279.3485283
- [67] D. Schneider, A. Otte, A. S. Kublin, A. Martschenko, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Accuracy of Commodity Finger Tracking Systems for Virtual Reality Head-Mounted Displays. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 804–805. IEEE, IEEE, 2020. doi: 10.1145/3485279.3485283
- [68] L. Sidenmark, C. Clarke, X. Zhang, J. Phu, and H. Gellersen. Outline Pursuits: Gaze-Assisted Selection of Occluded Objects in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2020. doi: 10.1145/3313831.3376438
- [69] P. Song, W. B. Goh, W. Hutama, C.-W. Fu, and X. Liu. A Handle Bar Metaphor for Virtual Object Manipulation With Mid-Air Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1297–1306, 2012. doi: 10.1145/2207676.2208585
- [70] L. Vanacken, T. Grossman, and K. Coninx. Exploring the Effects of Environment Density and Target Visibility on Object Selection in 3D Virtual Environments. In *2007 IEEE Symposium on 3D User Interfaces*. IEEE, IEEE, 2007. doi: 10.1145/3290605.3300331
- [71] R. Wang, S. Paris, and J. Popović. 6D Hands: Markerless Hand-Tracking for Computer Aided Design. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, pp. 549–558, 2011. doi: 10.1145/2047196.2047269
- [72] C. Wilkes and D. A. Bowman. Advantages of Velocity-Based Scaling for Distant 3D Manipulation. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*, VRST ’08, p. 23–29. Association for Computing Machinery, New York, NY, USA, 2008. doi: 10.1145/1450579.1450585
- [73] A. D. Wilson. Robust Computer Vision-Based Detection of Pinching for One and Two-Handed Gesture Input. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, UIST ’06, p. 255–258. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1166253.1166292
- [74] S. Winkler and D. Min. Stereo/multiview Picture Quality: Overview and Recent Advances. *Signal Processing: Image Communication*, 28(10):1358–1373, 2013. doi: 10.1016/j.image.2013.07.008
- [75] A. Worden, N. Walker, K. Bharat, and S. Hudson. Making Computers Easier for Older Adults to Use: Area Cursors and Sticky Icons. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 266–271, 1997. doi: 10.1145/258549.258724
- [76] A. Worden, N. Walker, K. Bharat, and S. Hudson. Making Computers Easier for Older Adults to Use: Area Cursors and Sticky Icons. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 266–271, 1997. doi: 10.1145/258549.258724
- [77] A. Zenner, H. M. Kriegler, and A. Krüger. Hart-The Virtual Reality Hand Redirection Toolkit. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2021. doi: 10.1145/3411763.3451814
- [78] A. Zenner and A. Krüger. Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 47–55. IEEE, IEEE, 2019. doi: 10.1007/978
- [79] A. Zenner, K. Ullmann, and A. Krüger. Combining Dynamic Passive Haptics and Haptic Retargeting for Enhanced Haptic Feedback in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2627–2637, 2021. doi: 10.1109/TVCG.2021.3067777
- [80] S. Zhai, W. Buxton, and P. Milgram. The “silk Cursor” Investigating Transparency for 3d Target Acquisition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 459–464, 1994. doi: 10.1145/191666.191822
- [81] S. Zhai, S. Conversy, M. Beaudouin-Lafon, and Y. Guiard. Human On-Line Response to Target Expansion. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 177–184, 2003. doi: 10.1145/642611.642644